	Type	L #	Hits	Search Text	DBs
1	BRS	L1	29676	(hash$3 AND (index OR indices OR key$3) AND (compare$ OR match$3))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
2	BRS	L2	800	((hash$3 AND (index OR indices OR key$3) AND (compar$3 OR match$3) AND (MAC NEAR3 Address$3)) AND pointer$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
3	BRS	L3	2010	(hash$3 AND (index OR indices OR key$3) AND (compar$3 OR match$3) AND (MAC NEAR3 Address$3))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB

	Type	L #	Hits	Search Text	DBs
4	BRS	L4	88	(search$3 or retriev$3 or register$3) with (hash$3) with index$3 and (memory adj2 table$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
5	BRS	L5	735	(search$3 or retriev$3 or register$3) with (hash$3) with index$3 and (memory same table$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
6	BRS	L6	2050	1 and (memory adj3 table$1) and hash$3 and (index$3 or entr$3 or key$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB

	Type	L #	Hits	Search Text	DBs
7	BRS	L7	376	(search$3 or retriev$3 or register$3) and (hash$3) and index$3 and (memory adj2 table$1) and (collision$3 or conflict$3 or synonym$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
8	BRS	L8	227	(Mori-Masaya or Takatsu-Yoshihisa or Watanabe-Shinpei or Sunaga-Toshio).in.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
9	BRS	L9	4	1 and 8	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB

	Type	L #	Hits	Search Text	DBs
10	BRS	L10	1065	(hash$3 AND (index OR indices OR key$3) AND (compare$ OR match$3)) and (memory adj2 table$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
11	BRS	L11	97	370/392,389,351,352,395,400,474.ccls. and 10	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
12	BRS	L12	688	(hash$3 AND (index OR indices OR key$3) AND (compare$ OR match$3)) and ((memory adj2 table$1) and pointer$1)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB

	Type	L #	Hits	Search Text	DBs
13	BRS	L14	17	12 and 709/238,243.ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
14	BRS	L15	32	12 and 711/216,220.ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
15	BRS	L16	16732	hash$3 AND (index OR indices OR key$3 OR indexes OR entry OR entries) AND (pointer$1 OR indirect$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB

	Type	L #	Hits	Search Text	DBs
16	BRS	L17	1704	hash$3 AND (index OR indices OR key$3 OR indexes OR entry OR entries) AND (memory same (pointer$1 and table$1) and indirect$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
17	BRS	L18	118	17 and 707/1-10.ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
18	BRS	L19	1439	hash$3 AND (index OR indices OR key$3 OR indexes OR entry OR entries) AND (memory same (addresses and pointer$1 and table$1) and indirect$3)	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB

	Type	L #	Hits	Search Text	DBs
19	BRS	L20	87	17 and 707/1,3,10.ccls.	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB
20	BRS	L21	537	(hash$3 AND (index OR indices OR key$3) AND (compare$ OR match$3)) and (memory and (addresses and table$1 pointer$1))	US-PGPUB; USPAT; USOCR; EPO; JPO; DERWENT; IBM_TDB

PORTAL
USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
+register* +hash* +index* +memory* +table*
SEARCH
THE ACM DIGITAL LIBRARY
Feedback

+register* +hash* +index* +memory* +table*
Terms used: **register hash index memory table**

Found **456** of **239,726**

Sort results by relevance | Save results to a Binder | Refine these results with Advanced Search

Display results expanded form | Open results in a new window | Try this search in The ACM Guide

Results 1 - 20 of 456 Result page: **1** 2 3 4 5 6 7 8 9 10 next >>

1 Scalable Hardware Memory Disambiguation for High ILP Processors
Simha Sethumadhavan, Rajagopalan Desikan, Doug Burger, Charles R. Moore, Stephen W. Keckler
December 2003 **MICRO 36:** Proceedings of the 36th annual IEEE/ACM International Symposium on Microarchitecture
Publisher: IEEE Computer Society
Full text available: pdf(220.70 KB) Additional Information: full citation, abstract, cited by, index terms

This paper describes several methods for improving thescalability of memory disambiguation hardware for futurehigh ILP processors. As the number of in-flight instructionsgrows with issue width and pipeline depth, the load/storequeues (LSQ) threaten to ...

2 Unbounded page-based transactional memory
Weihaw Chuang, Satish Narayanasamy, Ganesh Venkatesh, Jack Sampson, Michael Van Biesbrouck, Gilles Pokam, Brad Calder, Osvaldo Colavin
October 2006 **ACM SIGARCH Computer Architecture News**, Volume 34 Issue 5
Publisher: ACM
Full text available: pdf(242.68 KB) Additional Information: full citation, abstract, references, cited by, index terms

Exploiting thread level parallelism is paramount in the multicore era. Transactions enable programmers to expose such parallelism by greatly simplifying the multi-threaded programming model. Virtualized transactions (unbounded in space and time) are ...

Keywords: concurrency, parallel programming, transactional memory, transactions, virtual memory

3 Spatial Memory Streaming
Stephen Somogyi, Thomas F. Wenisch, Anastassia Ailamaki, Babak Falsafi, Andreas Moshovos
June 2006 **ISCA '06:** Proceedings of the 33rd annual international symposium on Computer Architecture
Publisher: IEEE Computer Society
Full text available: pdf(564.34 KB) Additional Information: full citation, abstract, references, cited by, index terms

Prior research indicates that there is much spatial variation in applications'

Ads by Google

Download PDF Converter
Convert Document & Image formats into PDF. Fast Download Guaranteed!
PDFConverter.PDF-forma

Information Management
Trust SAS to manage your data & information- Get Info Now!
www.SAS.com

Leading DSS & Analytic
Consulting for Health Plans. Download a Free Information Guide!
www.Medstat.com/Decisic

Paperless Meetings, ARS
Any web-based device can access surveys, Q&A, paperless meetings!
www.visiontree.com

memory access patterns. Modern memory systems, however, use small fixed-size cache blocks and as such cannot exploit the variation. Increasing the block size would not only ...

4 Exploiting frequent field values in java objects for reducing heap memory requirements

Guangyu Chen, Mahmut Kandemir, Mary J. Irwin
June 2005 **VEE '05:** Proceedings of the 1st ACM/USENIX international conference on Virtual execution environments
Publisher: ACM

Full text available: pdf(635.47 KB) Additional Information: full citation, abstract, references, cited by, index terms

The capabilities of applications executing on embedded and mobile devices are strongly influenced by memory size limitations. In fact, memory limitations are one of the main reasons that applications run slowly or even crash in embedded/mobile devices. ...

Keywords: Java virtual machine, frequent field value, garbage collection, heap

5 Unbounded page-based transactional memory

Weihaw Chuang, Satish Narayanasamy, Ganesh Venkatesh, Jack Sampson, Michael Van Biesbrouck, Gilles Pokam, Brad Calder, Osvaldo Colavin
November 2006 **ACM SIGPLAN Notices**, Volume 41 Issue 11
Publisher: ACM

Full text available: pdf(242.68 KB) Additional Information: full citation, abstract, references, cited by, index terms

Exploiting thread level parallelism is paramount in the multicore era. Transactions enable programmers to expose such parallelism by greatly simplifying the multi-threaded programming model. Virtualized transactions (unbounded in space and time) are ...

Keywords: concurrency, parallel programming, transactional memory, transactions, virtual memory

6 Memory forwarding: enabling aggressive layout optimizations by guaranteeing the safety of data relocation

Chi-Keung Luk, Todd C. Mowry
May 1999 **ACM SIGARCH Computer Architecture News**, Volume 27 Issue 2
Publisher: ACM

Full text available: pdf(196.77 KB) Publisher Site Additional Information: full citation, abstract, references, cited by, index terms

By optimizing data layout at run-time, we can potentially enhance the performance of caches by actively creating spatial locality, facilitating prefetching, and avoiding cache conflicts and false sharing. Unfortunately, it is extremely difficult to guarantee ...

7 Architectural Support for Software Transactional Memory

Bratin Saha, Ali-Reza Adl-Tabatabai, Quinn Jacobson
December 2006 **MICRO 39:** Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture
Publisher: IEEE Computer Society

Full text available: pdf(325.24 KB) Additional Information: full citation, abstract, references, cited by, index terms

Transactional memory provides a concurrency control mechanism that avoids many of the pitfalls of lock-based synchronization. Researchers have proposed several different implementations of transactional memory, broadly classified into software transactional ...

8 Informing memory operations: memory performance feedback mechanisms and their applications

Mark Horowitz, Margaret Martonosi, Todd C. Mowry, Michael D. Smith
May 1998 **ACM Transactions on Computer Systems (TOCS)**, Volume 16 Issue 2
Publisher: ACM

Full text available: pdf(344.74 KB) Additional Information: full citation, abstract, references, cited by, index terms, review

Memory latency is an important bottleneck in system performance that cannot be adequately solved by hardware alone. Several promising software techniques have been shown to address this problem successfully in specific situations. However, the generality ...

Keywords: cache miss notification, memory latency, processor architecture

9 Address-Indexed Memory Disambiguation and Store-to-Load Forwarding

Sam S. Stone, Kevin M. Woley, Matthew I. Frank
November 2005 **MICRO 38:** Proceedings of the 38th annual IEEE/ACM International Symposium on Microarchitecture
Publisher: IEEE Computer Society

Full text available: pdf(301.30 KB) Publisher Site Additional Information: full citation, abstract, references, cited by, index terms

This paper describes a scalable, low-complexity alternative to the conventional load/store queue (LSQ) for superscalar processors that execute load and store instructions speculatively and out-of-order prior to resolving their dependences. Whereas the ...

10 Spatial Memory Streaming

Stephen Somogyi, Thomas F. Wenisch, Anastassia Ailamaki, Babak Falsafi, Andreas Moshovos
May 2006 **ACM SIGARCH Computer Architecture News**, Volume 34 Issue 2
Publisher: ACM

Full text available: pdf(564.34 KB) Additional Information: full citation, abstract, references, cited by, index terms

Prior research indicates that there is much spatial variation in applications' memory access patterns. Modern memory systems, however, use small fixed-size cache blocks and as such cannot exploit the variation. Increasing the block size would not only ...

11 Fighting the memory wall with assisted execution

Michel Dubois
April 2004 **CF '04:** Proceedings of the 1st conference on Computing frontiers
Publisher: ACM

Full text available: pdf(231.18 KB) Additional Information: full citation, abstract, references, cited by, index terms

Assisted execution is a form of simultaneous multithreading in which a set of auxiliary "assistant" threads, called *nanothreads*, is attached to each thread of an application. Nanothreads are lightweight threads which run on the same processor ...

Keywords: cache memories, latency tolerance, prefetching, simultaneous multithreading, superscalar processors

12 Virtualizing Transactional Memory

Ravi Rajwar, Maurice Herlihy, Konrad Lai
June 2005 **ISCA '05:** Proceedings of the 32nd annual international symposium on Computer Architecture
Publisher: IEEE Computer Society

Full text available: pdf(199.77 KB) Additional Information: full citation, abstract, cited by, index terms

Writing concurrent programs is difficult because of the complexity of ensuring proper synchronization. Conventional lock-based synchronization suffers from wellknown limitations, so researchers have considered non-blocking transactions as an alternative. ...

13 Hybrid transactional memory

Sanjeev Kumar, Michael Chu, Christopher J. Hughes, Partha Kundu, Anthony Nguyen
March 2006 **PPoPP '06:** Proceedings of the eleventh ACM SIGPLAN symposium on Principles and practice of parallel programming
Publisher: ACM

Full text available: pdf(782.57 KB) Additional Information: full citation, abstract, references, cited by, index terms

High performance parallel programs are currently difficult to write and debug. One major source of difficulty is protecting concurrent accesses to shared data with an appropriate synchronization mechanism. Locks are the most common mechanism but they ...

Keywords: architecture support, nonblocking, transactional memory, transactions

14 High performance dynamic lock-free hash tables and list-based sets

Maged M. Michael
August 2002 **SPAA '02:** Proceedings of the fourteenth annual ACM symposium on Parallel algorithms and architectures
Publisher: ACM

Full text available: pdf(238.11 KB) Additional Information: full citation, abstract, references, cited by, index terms

Lock-free (non-blocking) shared data structures promise more robust performance and reliability than conventional lock-based implementations.

However, all prior lock-free algorithms for sets and hash tables suffer from serious drawbacks that prevent ...

15 Memory-efficient content filtering hardware for high-speed intrusion detection systems

Sungwon Yi, Byoung-koo Kim, Jintae Oh, Jongsoo Jang, George Kesidis, Chita R. Das

March 2007 **SAC '07:** Proceedings of the 2007 ACM symposium on Applied computing

Publisher: ACM

Full text available: pdf(161.42 KB) Additional Information: full citation, abstract, references, index terms

Content filtering-based Intrusion Detection Systems have been widely deployed in enterprise networks, and have become a standard measure to protect networks and network users from cyber attacks. Although several solutions have been proposed recently, ...

Keywords: FPGA, content filtering, deep packet instpection, intrusion detection systems, network security, pattern matching

16 Memory Protection through Dynamic Access Control

Kun Zhang, Tao Zhang, Santosh Pande

December 2006 **MICRO 39:** Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture

Publisher: IEEE Computer Society

Full text available: pdf(274.63 KB) Additional Information: full citation, abstract, references, index terms

Current anomaly detection schemes focus on control flow monitoring. Recently, Chen et al. [2] discovered that a large category of attacks tamper program data but do not alter control flows. These attacks are not only realistic, but are also as important ...

17 Selection conditions in main memory

Kenneth A. Ross

March 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 1

Publisher: ACM

Full text available: pdf(296.54 KB) Additional Information: full citation, abstract, references, cited by, index terms

We consider the fundamental operation of applying a compound filtering condition to a set of records. With large main memories available cheaply, systems may choose to keep the data entirely in main memory, in order to improve query and/or update performance.The ...

Keywords: Branch misprediction

18 Tolerating memory latency through push prefetching for pointer-intensive applications

Chia-Lin Yang, Alvin R. Lebeck, Hung-Wei Tseng, Chien-Hao Lee

December 2004 **ACM Transactions on Architecture and Code Optimization (TACO)**, Volume 1 Issue 4

Publisher: ACM

Full text available: pdf(590.24 KB) Additional Information: full citation, abstract, references, cited by, index terms

Prefetching is often used to overlap memory latency with computation for array-based applications. However, prefetching for pointer-intensive applications remains a challenge because of the irregular memory access pattern and pointer-chasing problem. ...

Keywords: Prefetch, linked data structures, memory hierarchy, pointer-chasing

19 A study of Erlang ETS table implementations and performance

Scott Lystig Fritchie

August 2003 **ERLANG '03:** Proceedings of the 2003 ACM SIGPLAN workshop on Erlang

Publisher: ACM

Full text available: pdf(232.17 KB) Additional Information: full citation, abstract, references, index terms

The viability of implementing an in-memory database, Erlang ETS, using a relatively-new data structure, called a Judy array, was studied by comparing the performance of ETS tables based on four data structures: AVL balanced binary trees, B-trees, resizable ...

Keywords: AVL tree, B-tree, Erlang, Judy array, hash table, in-memory database

20 An integrated hardware-software approach to flexible transactional memory

Arrvindh Shriraman, Michael F. Spear, Hemayet Hossain, Virendra J. Marathe, Sandhya Dwarkadas, Michael L. Scott

June 2007 **ISCA '07:** Proceedings of the 34th annual international symposium on Computer architecture

Publisher: ACM

Full text available: pdf(637.98 KB) Additional Information: full citation, abstract, references, index terms

There has been considerable recent interest in both hardware andsoftware transactional memory (TM). We present an intermediateapproach, in which hardware serves to accelerate a TM implementation controlled fundamentally by software. Specifically, we ...

Keywords: RSTM, cache coherence, multiprocessors, transactional memory

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2008 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

PORTAL
USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
+register* +hash* +index* +memory* +table*
SEARCH
THE ACM DIGITAL LIBRARY
Feedback

+register* +hash* +index* +memory* +table*
Terms used: **register hash index memory table**

Found **456** of **239,726**

Sort results by relevance | Save results to a Binder | Refine these results with Advanced Search

Display results expanded form | Open results in a new window | Try this search in The ACM Guide

Results 1 - 20 of 456 Result page: **1** 2 3 4 5 6 7 8 9 10 next >>

1 Scalable Hardware Memory Disambiguation for High ILP Processors
Simha Sethumadhavan, Rajagopalan Desikan, Doug Burger, Charles R. Moore, Stephen W. Keckler
December 2003 **MICRO 36:** Proceedings of the 36th annual IEEE/ACM International Symposium on Microarchitecture
Publisher: IEEE Computer Society
Full text available: pdf(220.70 KB) Additional Information: full citation, abstract, cited by, index terms

This paper describes several methods for improving thescalability of memory disambiguation hardware for futurehigh ILP processors. As the number of in-flight instructionsgrows with issue width and pipeline depth, the load/storequeues (LSQ) threaten to ...

2 Unbounded page-based transactional memory
Weihaw Chuang, Satish Narayanasamy, Ganesh Venkatesh, Jack Sampson, Michael Van Biesbrouck, Gilles Pokam, Brad Calder, Osvaldo Colavin
October 2006 **ACM SIGARCH Computer Architecture News**, Volume 34 Issue 5
Publisher: ACM
Full text available: pdf(242.68 KB) Additional Information: full citation, abstract, references, cited by, index terms

Exploiting thread level parallelism is paramount in the multicore era. Transactions enable programmers to expose such parallelism by greatly simplifying the multi-threaded programming model. Virtualized transactions (unbounded in space and time) are ...

Keywords: concurrency, parallel programming, transactional memory, transactions, virtual memory

3 Spatial Memory Streaming
Stephen Somogyi, Thomas F. Wenisch, Anastassia Ailamaki, Babak Falsafi, Andreas Moshovos
June 2006 **ISCA '06:** Proceedings of the 33rd annual international symposium on Computer Architecture
Publisher: IEEE Computer Society
Full text available: pdf(564.34 KB) Additional Information: full citation, abstract, references, cited by, index terms

Prior research indicates that there is much spatial variation in applications'

Ads by Google

Download PDF Converter
Convert Document & Image formats into PDF. Fast Download Guaranteed!
PDFConverter.PDF-forma

Information Management
Trust SAS to manage your data & information- Get Info Now!
www.SAS.com

Leading DSS & Analytic
Consulting for Health Plans. Download a Free Information Guide!
www.Medstat.com/Decisic

Paperless Meetings, ARS
Any web-based device can access surveys, Q&A, paperless meetings!
www.visiontree.com

memory access patterns. Modern memory systems, however, use small fixed-size cache blocks and as such cannot exploit the variation. Increasing the block size would not only ...

4 Exploiting frequent field values in java objects for reducing heap memory requirements

Guangyu Chen, Mahmut Kandemir, Mary J. Irwin
June 2005 **VEE '05:** Proceedings of the 1st ACM/USENIX international conference on Virtual execution environments
Publisher: ACM

Full text available: pdf(635.47 KB) Additional Information: full citation, abstract, references, cited by, index terms

The capabilities of applications executing on embedded and mobile devices are strongly influenced by memory size limitations. In fact, memory limitations are one of the main reasons that applications run slowly or even crash in embedded/mobile devices. ...

Keywords: Java virtual machine, frequent field value, garbage collection, heap

5 Unbounded page-based transactional memory

Weihaw Chuang, Satish Narayanasamy, Ganesh Venkatesh, Jack Sampson, Michael Van Biesbrouck, Gilles Pokam, Brad Calder, Osvaldo Colavin
November 2006 **ACM SIGPLAN Notices**, Volume 41 Issue 11
Publisher: ACM

Full text available: pdf(242.68 KB) Additional Information: full citation, abstract, references, cited by, index terms

Exploiting thread level parallelism is paramount in the multicore era. Transactions enable programmers to expose such parallelism by greatly simplifying the multi-threaded programming model. Virtualized transactions (unbounded in space and time) are ...

Keywords: concurrency, parallel programming, transactional memory, transactions, virtual memory

6 Memory forwarding: enabling aggressive layout optimizations by guaranteeing the safety of data relocation

Chi-Keung Luk, Todd C. Mowry
May 1999 **ACM SIGARCH Computer Architecture News**, Volume 27 Issue 2
Publisher: ACM

Full text available: pdf(196.77 KB) Publisher Site Additional Information: full citation, abstract, references, cited by, index terms

By optimizing data layout at run-time, we can potentially enhance the performance of caches by actively creating spatial locality, facilitating prefetching, and avoiding cache conflicts and false sharing. Unfortunately, it is extremely difficult to guarantee ...

7 Architectural Support for Software Transactional Memory

Bratin Saha, Ali-Reza Adl-Tabatabai, Quinn Jacobson
December 2006 **MICRO 39:** Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture
Publisher: IEEE Computer Society

Full text available: pdf(325.24 KB) Additional Information: full citation, abstract, references, cited by, index terms

Transactional memory provides a concurrency control mechanism that avoids many of the pitfalls of lock-based synchronization. Researchers have proposed several different implementations of transactional memory, broadly classified into software transactional ...

8 Informing memory operations: memory performance feedback mechanisms and their applications

Mark Horowitz, Margaret Martonosi, Todd C. Mowry, Michael D. Smith
May 1998 **ACM Transactions on Computer Systems (TOCS)**, Volume 16 Issue 2
Publisher: ACM

Full text available: pdf(344.74 KB) Additional Information: full citation, abstract, references, cited by, index terms, review

Memory latency is an important bottleneck in system performance that cannot be adequately solved by hardware alone. Several promising software techniques have been shown to address this problem successfully in specific situations. However, the generality ...

Keywords: cache miss notification, memory latency, processor architecture

9 Address-Indexed Memory Disambiguation and Store-to-Load Forwarding

Sam S. Stone, Kevin M. Woley, Matthew I. Frank
November 2005 **MICRO 38:** Proceedings of the 38th annual IEEE/ACM International Symposium on Microarchitecture
Publisher: IEEE Computer Society

Full text available: pdf(301.30 KB) Publisher Site Additional Information: full citation, abstract, references, cited by, index terms

This paper describes a scalable, low-complexity alternative to the conventional load/store queue (LSQ) for superscalar processors that execute load and store instructions speculatively and out-of-order prior to resolving their dependences. Whereas the ...

10 Spatial Memory Streaming

Stephen Somogyi, Thomas F. Wenisch, Anastassia Ailamaki, Babak Falsafi, Andreas Moshovos
May 2006 **ACM SIGARCH Computer Architecture News**, Volume 34 Issue 2
Publisher: ACM

Full text available: pdf(564.34 KB) Additional Information: full citation, abstract, references, cited by, index terms

Prior research indicates that there is much spatial variation in applications' memory access patterns. Modern memory systems, however, use small fixed-size cache blocks and as such cannot exploit the variation. Increasing the block size would not only ...

11 Fighting the memory wall with assisted execution

Michel Dubois
April 2004 **CF '04:** Proceedings of the 1st conference on Computing frontiers
Publisher: ACM

Full text available: pdf(231.18 KB) Additional Information: full citation, abstract, references, cited by, index terms

Assisted execution is a form of simultaneous multithreading in which a set of auxiliary "assistant" threads, called *nanothreads*, is attached to each thread of an application. Nanothreads are lightweight threads which run on the same processor ...

Keywords: cache memories, latency tolerance, prefetching, simultaneous multithreading, superscalar processors

12 Virtualizing Transactional Memory

Ravi Rajwar, Maurice Herlihy, Konrad Lai
June 2005 **ISCA '05:** Proceedings of the 32nd annual international symposium on Computer Architecture
Publisher: IEEE Computer Society

Full text available: pdf(199.77 KB) Additional Information: full citation, abstract, cited by, index terms

Writing concurrent programs is difficult because of the complexity of ensuring proper synchronization. Conventional lock-based synchronization suffers from wellknown limitations, so researchers have considered non-blocking transactions as an alternative. ...

13 Hybrid transactional memory

Sanjeev Kumar, Michael Chu, Christopher J. Hughes, Partha Kundu, Anthony Nguyen
March 2006 **PPoPP '06:** Proceedings of the eleventh ACM SIGPLAN symposium on Principles and practice of parallel programming
Publisher: ACM

Full text available: pdf(782.57 KB) Additional Information: full citation, abstract, references, cited by, index terms

High performance parallel programs are currently difficult to write and debug. One major source of difficulty is protecting concurrent accesses to shared data with an appropriate synchronization mechanism. Locks are the most common mechanism but they ...

Keywords: architecture support, nonblocking, transactional memory, transactions

14 High performance dynamic lock-free hash tables and list-based sets

Maged M. Michael
August 2002 **SPAA '02:** Proceedings of the fourteenth annual ACM symposium on Parallel algorithms and architectures
Publisher: ACM

Full text available: pdf(238.11 KB) Additional Information: full citation, abstract, references, cited by, index terms

Lock-free (non-blocking) shared data structures promise more robust performance and reliability than conventional lock-based implementations.

However, all prior lock-free algorithms for sets and hash tables suffer from serious drawbacks that prevent ...

15 Memory-efficient content filtering hardware for high-speed intrusion detection systems

Sungwon Yi, Byoung-koo Kim, Jintae Oh, Jongsoo Jang, George Kesidis, Chita R. Das

March 2007 **SAC '07:** Proceedings of the 2007 ACM symposium on Applied computing

Publisher: ACM

Full text available: pdf(161.42 KB) Additional Information: full citation, abstract, references, index terms

Content filtering-based Intrusion Detection Systems have been widely deployed in enterprise networks, and have become a standard measure to protect networks and network users from cyber attacks. Although several solutions have been proposed recently, ...

Keywords: FPGA, content filtering, deep packet instpection, intrusion detection systems, network security, pattern matching

16 Memory Protection through Dynamic Access Control

Kun Zhang, Tao Zhang, Santosh Pande

December 2006 **MICRO 39:** Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture

Publisher: IEEE Computer Society

Full text available: pdf(274.63 KB) Additional Information: full citation, abstract, references, index terms

Current anomaly detection schemes focus on control flow monitoring. Recently, Chen et al. [2] discovered that a large category of attacks tamper program data but do not alter control flows. These attacks are not only realistic, but are also as important ...

17 Selection conditions in main memory

Kenneth A. Ross

March 2004 **ACM Transactions on Database Systems (TODS)**, Volume 29 Issue 1

Publisher: ACM

Full text available: pdf(296.54 KB) Additional Information: full citation, abstract, references, cited by, index terms

We consider the fundamental operation of applying a compound filtering condition to a set of records. With large main memories available cheaply, systems may choose to keep the data entirely in main memory, in order to improve query and/or update performance.The ...

Keywords: Branch misprediction

18 Tolerating memory latency through push prefetching for pointer-intensive applications

Chia-Lin Yang, Alvin R. Lebeck, Hung-Wei Tseng, Chien-Hao Lee

December 2004 **ACM Transactions on Architecture and Code Optimization (TACO)**, Volume 1 Issue 4

Publisher: ACM

Full text available: pdf(590.24 KB) Additional Information: full citation, abstract, references, cited by, index terms

Prefetching is often used to overlap memory latency with computation for array-based applications. However, prefetching for pointer-intensive applications remains a challenge because of the irregular memory access pattern and pointer-chasing problem. ...

Keywords: Prefetch, linked data structures, memory hierarchy, pointer-chasing

19 A study of Erlang ETS table implementations and performance

Scott Lystig Fritchie

August 2003 **ERLANG '03:** Proceedings of the 2003 ACM SIGPLAN workshop on Erlang

Publisher: ACM

Full text available: pdf(232.17 KB) Additional Information: full citation, abstract, references, index terms

The viability of implementing an in-memory database, Erlang ETS, using a relatively-new data structure, called a Judy array, was studied by comparing the performance of ETS tables based on four data structures: AVL balanced binary trees, B-trees, resizable ...

Keywords: AVL tree, B-tree, Erlang, Judy array, hash table, in-memory database

20 An integrated hardware-software approach to flexible transactional memory

Arrvindh Shriraman, Michael F. Spear, Hemayet Hossain, Virendra J. Marathe, Sandhya Dwarkadas, Michael L. Scott

June 2007 **ISCA '07:** Proceedings of the 34th annual international symposium on Computer architecture

Publisher: ACM

Full text available: pdf(637.98 KB) Additional Information: full citation, abstract, references, index terms

There has been considerable recent interest in both hardware andsoftware transactional memory (TM). We present an intermediateapproach, in which hardware serves to accelerate a TM implementation controlled fundamentally by software. Specifically, we ...

Keywords: RSTM, cache coherence, multiprocessors, transactional memory

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2008 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

PORTAL
USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
+MAC* +hash* +index* +memory* +table*
SEARCH
THE ACM DIGITAL LIBRARY
Feedback

+MAC* +hash* +index* +memory* +table*
Terms used: **MAC hash index memory table**

Found **153** of **239,726**

Sort results by: relevance | Save results to a Binder
Display results: expanded form | Open results in a new window

Refine these results with Advanced Search
Try this search in The ACM Guide

Results 1 - 20 of 153 Result page: **1** 2 3 4 5 6 7 8 next >>

1 Memory forwarding: enabling aggressive layout optimizations by guaranteeing the safety of data relocation
Chi-Keung Luk, Todd C. Mowry
May 1999 **ACM SIGARCH Computer Architecture News**, Volume 27 Issue 2
Publisher: ACM

Full text available: pdf(196.77 KB) Publisher Site
Additional Information: full citation, abstract, references, cited by, index terms

By optimizing data layout at run-time, we can potentially enhance the performance of caches by actively creating spatial locality, facilitating prefetching, and avoiding cache conflicts and false sharing. Unfortunately, it is extremely difficult to guarantee ...

2 Attested append-only memory: making adversaries stick to their word
Byung-Gon Chun, Petros Maniatis, Scott Shenker, John Kubiatowicz
October 2007 **SOSP '07:** Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles
Publisher: ACM

Full text available: pdf(361.31 KB) Additional Information: full citation, abstract, references, index terms

Researchers have made great strides in improving the fault tolerance of both centralized and replicated systems against arbitrary (Byzantine) faults. However, there are hard limits to how much can be done with entirely untrusted components; for example, ...

Keywords: attested append-only memory, byzantine-fault tolerance, equivocation, replicated state machines, shared storage

3 Efficient use of memory bandwidth to improve network processor throughput
Jahangir Hasan, Satish Chandra, T. N. Vijaykumar
June 2003 **ISCA '03:** Proceedings of the 30th annual international symposium on Computer architecture
Publisher: ACM

Additional Information:

Ads by Google

Download PDF Converter
Convert Document & Image formats into PDF. Fast Download Guaranteed!
PDFConverter.PDF-forma

Information Management
Trust SAS to manage your data & information- Get Info Now!
www.SAS.com

Leading DSS & Analytic
Consulting for Health Plans. Download a Free Information Guide!
www.Medstat.com/Decisic

Paperless Meetings, ARS
Any web-based device can access surveys, Q&A, paperless meetings!
www.visiontree.com

Full text available: pdf(184.83 KB) full citation, abstract, references, cited by

We consider the efficiency of packet buffers used in packet switches built using network processors (NPs). Packet buffers are typically implemented using DRAM, which provides plentiful buffering at a reasonable cost. The problem we address is that a ...

4 Virtually Pipelined Network Memory
Banit Agrawal, Timothy Sherwood
December 2006 **MICRO 39:** Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture
Publisher: IEEE Computer Society
Full text available: pdf(256.77 KB) Additional Information: full citation, abstract, references, index terms

We introduce virtually-pipelined memory, an architectural technique that efficiently supports high-bandwidth, uniform latency memory accesses, and high-confidence throughput even under adversarial conditions. We apply this technique to the network processing ...

5 Architectural Support for High Speed Protection of Memory Integrity and Confidentiality in Multiprocessor Systems
Weidong Shi, Hsien-Hsin S. Lee, Mrinmoy Ghosh, Chenghuai Lu
September 2004 **PACT '04:** Proceedings of the 13th International Conference on Parallel Architectures and Compilation Techniques
Publisher: IEEE Computer Society
Full text available: pdf(255.33 KB) Additional Information: full citation, abstract, cited by

Recently there is a growing effort in both the architecture and the security community to create a hardware solution for authenticating system memory.As shown in the previous work, hardware-based memory authentication will become a vital component for ...

6 Efficient use of memory bandwidth to improve network processor throughput
Jahangir Hasan, Satish Chandra, T. N. Vijaykumar
May 2003 **ACM SIGARCH Computer Architecture News**, Volume 31 Issue 2
Publisher: ACM
Full text available: pdf(184.83 KB) Additional Information: full citation, abstract, references, cited by

We consider the efficiency of packet buffers used in packet switches built using network processors (NPs). Packet buffers are typically implemented using DRAM, which provides plentiful buffering at a reasonable cost. The problem we address is that a ...

7 Memory forwarding: enabling aggressive layout optimizations by guaranteeing the safety of data relocation
Chi-Keung Luk, Todd C. Mowry
May 1999 **ISCA '99:** Proceedings of the 26th annual international

symposium on Computer architecture
Publisher: IEEE Computer Society

Additional Information: full citation, abstract, references, cited by, index terms

Full text available: pdf(196.77 KB) Publisher Site

By optimizing data layout at run-time, we can potentially enhance the performance of caches by actively creating spatial locality, facilitating prefetching, and avoiding cache conflicts and false sharing. Unfortunately, it is extremely difficult to guarantee ...

8 Accelerating sparse matrix computations via data compression

Jeremiah Willcock, Andrew Lumsdaine
June 2006 **ICS '06:** Proceedings of the 20th annual international conference on Supercomputing
Publisher: ACM

Full text available: pdf(652.09 KB) Additional Information: full citation, abstract, references, index terms

Sparse matrix computations are important for many scientific computations, with matrix-vector multiplication being a fundamental operation for modern iterative algorithms. For large sparse matrices, the primary performance limitation on matrix-vector ...

Keywords: data compression, high-performance computing, memory bandwidth, sparse matrix

9 Trajectory sampling for direct traffic observation

N. G. Duffield, Matthias Grossglauser
June 2001 **IEEE/ACM Transactions on Networking (TON)**, Volume 9 Issue 3
Publisher: IEEE Press

Additional Information: full citation, abstract, references, cited by, index terms

Full text available: pdf(251.55 KB)

Traffic measurement is a critical component for the control and engine ering of communication networks. We argue that traffic measurement should make it possible to obtain the spatial flow of traffic through the domain, i.e., the paths followed by packets ...

Keywords: Hash functions, Internet traffic measurement, packet sampling, traffic engineering

10 Algorithms and data structures for flash memories

Eran Gal, Sivan Toledo
June 2005 **ACM Computing Surveys (CSUR)**, Volume 37 Issue 2
Publisher: ACM

Additional Information: full citation, abstract, references, cited by, index terms

Full text available: pdf(343.39 KB)

Flash memory is a type of electrically-erasable programmable read-only memory (EEPROM). Because flash memories are nonvolatile and relatively dense, they are now used to store files and other persistent objects in handheld computers, mobile phones, digital ...

Keywords: EEPROM memory, Flash memory, wear leveling

11 DejaView: a personal virtual computer recorder

Oren Laadan, Ricardo A. Baratto, Dan B. Phung, Shaya Potter, Jason Nieh
October 2007 **SOSP '07:** Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles
Publisher: ACM

Full text available: pdf(534.51 KB) Additional Information: full citation, abstract, references, index terms

As users interact with the world and their peers through their computers, it is becoming important to archive and later search the information that they have *viewed*. We present DejaView, a personal virtual computer recorder that provides a complete ...

Keywords: desktop search, virtualization

12 The V-Way Cache: Demand Based Associativity via Global Replacement

Moinuddin K. Qureshi, David Thompson, Yale N. Patt
May 2005 **ACM SIGARCH Computer Architecture News**, Volume 33 Issue 2
Publisher: ACM

Full text available: pdf(231.93 KB) Additional Information: full citation, abstract, cited by, index terms

As processor speeds increase and memory latency becomes more critical, intelligent design and management of secondary caches becomes increasingly important. The efficiency of current set-associative caches is reduced because programs exhibit a non-uniform ...

13 SpliceNP: a TCP splicer using a network processor

Li Zhao, Yan Luo, Laxmi Bhuyan, Ravi Iyer
October 2005 **ANCS '05:** Proceedings of the 2005 ACM symposium on Architecture for networking and communications systems
Publisher: ACM

Full text available: pdf(141.80 KB) Additional Information: full citation, abstract, references, cited by, index terms

TCP Splicing can be used in content-aware switches to tremendously reduce overall request latency. In order to reduce the processing latency further, we propose to offload the protocol processing onto network processors (NPs). An NP consists of a multithreaded ...

Keywords: TCP splicing, network processors

14 A history of Erlang

Joe Armstrong
June 2007 **HOPL III:** Proceedings of the third ACM SIGPLAN conference on History of programming languages
Publisher: ACM

Full text available: pdf(446.07 KB) Additional Information: full citation, appendices and supplements, abstract,

references, index terms

Erlang was designed for writing concurrent programs that "run forever." Erlang uses concurrent processes to structure the program. These processes have no shared memory and communicate by asynchronous message passing. Erlang processes are lightweight ...

15 Threats to privacy in the forensic analysis of database systems

Patrick Stahlberg, Gerome Miklau, Brian Neil Levine
June 2007 **SIGMOD '07:** Proceedings of the 2007 ACM SIGMOD international conference on Management of data
Publisher: ACM

Full text available: pdf(457.59 KB) Additional Information: full citation, abstract, references, index terms

The use of any modern computer system leaves unintended traces of expired data and remnants of users' past activities. In this paper, we investigate the unintended persistence of data stored in database systems. This data can be recovered by forensic ...

Keywords: forensics, privacy, transparency

16 Open multi-methods for c++

Peter Pirkelbauer, Yuriy Solodkyy, Bjarne Stroustrup
October 2007 **GPCE '07:** Proceedings of the 6th international conference on Generative programming and component engineering
Publisher: ACM

Full text available: pdf(427.67 KB) Additional Information: full citation, abstract, references, index terms

Multiple dispatch - the selection of a function to be invoked based on the dynamic type of two or more arguments - is a solution to several classical problems in object-oriented programming. Open multi-methods generalize multiple dispatch towards open-class ...

Keywords: C++, generic programming, multi-methods, multiple dispatch, object oriented programming, open-methods

17 Tree bitmap: hardware/software IP lookups with incremental updates

Will Eatherton, George Varghese, Zubin Dittia
April 2004 **ACM SIGCOMM Computer Communication Review**, Volume 34 Issue 2
Publisher: ACM

Full text available: pdf(189.39 KB) Additional Information: full citation, abstract, references, cited by

Even with the significant focus on IP address lookup in the published literature as well as focus on this market by commercial semiconductor vendors, there is still a challenge for router architects to find solutions that simultaneously meet 3 criteria: ...

18 Implementing an untrusted operating system on trusted hardware

David Lie, Chandramohan A. Thekkath, Mark Horowitz
December 2003 **ACM SIGOPS Operating Systems Review**, Volume 37 Issue

5

Publisher: ACM

Full text available: pdf(280.87 KB) Additional Information: full citation, abstract, references, cited by, index terms

Recently, there has been considerable interest in providing "trusted computing platforms" using hardware~---~TCPA and Palladium being the most publicly visible examples. In this paper we discuss our experience with building such a platform using a traditional ...

Keywords: XOM, XOMOS, untrusted operating systems

19 CAPTRA: coordinated packet traceback

Denh Sy, Lichun Bao

April 2006 **IPSN '06:** Proceedings of the fifth international conference on Information processing in sensor networks

Publisher: ACM

Full text available: pdf(307.50 KB) Additional Information: full citation, abstract, references, index terms

Network-based attacks can be either persistent or sporadic. Persistent attack flows can be relatively easy to trace by mechanisms such as probabilistic packet marking, traffic logging, data mining etc. Sporadic attacks are sometimes easily detected by ...

Keywords: bloom filter, packet traceback, wireless sensor networks

20 NPSE: A High Performance Network Packet Search Engine

Naresh Soni, Nick Richardson, Lun-Bin Huang, Suresh Rajgopal, George Vlantis

March 2003 **DATE '03: Proceedings of the conference on Design, Automation and Test in Europe: Designers' Forum - Volume 2**, Volume 2

Publisher: IEEE Computer Society

Full text available: pdf(118.36 KB) Publisher Site Additional Information: full citation, abstract, cited by, index terms

This paper describes the NPSE, a high-performance SRAM-based network packet search engine which has the primary application of supporting IPv4 and IPv6 forwarding. It is based on a high-speed hardware implementation of a tree-based storage and retrieval ...

Results 1 - 20 of 153 Result page: **1** 2 3 4 5 6 7 8 next >>

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2008 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

PORTAL
USPTO
Subscribe (Full Service) Register (Limited Service, Free) Login
Search: The ACM Digital Library The Guide
+register* +address* +first hash* +second hash*
SEARCH
THE ACM DIGITAL LIBRARY
Feedback

+register* +address* +first hash* +second hash*
Terms used: **register address first hash second hash**

Found **983** of **239,726**

Sort results by relevance | Save results to a Binder | Refine these results with Advanced Search

Display results expanded form | Open results in a new window | Try this search in The ACM Guide

Results 1 - 20 of 983 Result page: **1** 2 3 4 5 6 7 8 9 10 next >>

1 Non-uniform key distribution and address calculation sorting
Michael E. Tarter, Richard A. Kronmal
July 1966 **Communications of the ACM**, Volume 9 Issue 7
Publisher: ACM
Full text available: pdf(654.43 KB) Additional Information: full citation

2 First price sealed bid auction without auctioneers
Shuo Zheng, Luke McAven, Yi Mu
August 2007 **IWCMC '07:** Proceedings of the 2007 international conference on Wireless communications and mobile computing
Publisher: ACM
Full text available: pdf(129.27 KB) Additional Information: full citation, abstract, references, index terms

We propose two protocol variants for a first price sealed-bid auction, without using intermediatory auctioneers. One version achieves full privacy for the bidders and their bids, the other provides a form of verifiability, at the cost of some privacy. ...

Keywords: auction, bit-commitment, first price, sealed-bid

3 Address-Indexed Memory Disambiguation and Store-to-Load Forwarding
Sam S. Stone, Kevin M. Woley, Matthew I. Frank
November 2005 **MICRO 38:** Proceedings of the 38th annual IEEE/ACM International Symposium on Microarchitecture
Publisher: IEEE Computer Society
Full text available: pdf(301.30 KB) Publisher Site Additional Information: full citation, abstract, references, cited by, index terms

This paper describes a scalable, low-complexity alternative to the conventional load/store queue (LSQ) for superscalar processors that execute load and store instructions speculatively and out-of-order prior to resolving their dependences. Whereas the ...

4 A compact FPGA implementation of the hash function whirlpool

Ads by Google

Download PDF Converter
Convert Document & Image formats into PDF. Fast Download Guaranteed!
PDFConverter.PDF-forma

Information Management
Trust SAS to manage your data & information- Get Info Now!
www.SAS.com

Leading DSS & Analytic
Consulting for Health Plans. Download a Free Information Guide!
www.Medstat.com/Decisic

Paperless Meetings, ARS
Any web-based device can access surveys, Q&A, paperless meetings!
www.visiontree.com

Norbert Pramstaller, Christian Rechberger, Vincent Rijmen
February 2006 **FPGA '06:** Proceedings of the 2006 ACM/SIGDA 14th international symposium on Field programmable gate arrays
Publisher: ACM

Full text available: pdf(240.32 KB) Additional Information: full citation, abstract, references, index terms

Recent breakthroughs in cryptanalysis of standard hash functions like SHA-1 and MD5 raise the need for alternatives. A credible alternative to for instance SHA-1 or the SHA-2 family of hash functions is Whirlpool. Whirlpool is a hash function that has ...

Keywords: FPGA, compact hardware implementation, hash function, whirlpool

5 Repairing return address stack for buffer overflow protection
Yong-Joon Park, Gyungho Lee
April 2004 **CF '04:** Proceedings of the 1st conference on Computing frontiers
Publisher: ACM

Full text available: pdf(197.90 KB) Additional Information: full citation, abstract, references, cited by, index terms

Although many defense mechanisms against buffer overflow attacks have been proposed, buffer overflow vulnerability in software is still one of the most prevalent vulnerabilities exploited. This paper proposes a micro-architecture based defense mechanism ...

Keywords: buffer overflow, computer architecture, computer security, intrusion tolerance

6 Implementing aggregation and broadcast over Distributed Hash Tables
Ji Li, Karen Sollins, Dah-Yoh Lim
January 2005 **ACM SIGCOMM Computer Communication Review**, Volume 35 Issue 1
Publisher: ACM

Full text available: pdf(391.51 KB) Additional Information: full citation, abstract, references, cited by, index terms

Peer-to-peer (P2P) networks represent an effective way to share information, since there are no central points of failure or bottleneck. However, the flip side to the distributive nature of P2P networks is that it is not trivial to aggregate and broadcast ...

Keywords: aggregation, broadcast, distributed hash table, peer-to-peer, tree

7 High performance dynamic lock-free hash tables and list-based sets
Maged M. Michael
August 2002 **SPAA '02:** Proceedings of the fourteenth annual ACM symposium on Parallel algorithms and architectures
Publisher: ACM

Full text available: pdf(238.11 KB) Additional Information: full citation, abstract, references, cited by, index terms

Lock-free (non-blocking) shared data structures promise more robust performance and reliability than conventional lock-based implementations. However, all prior lock-free algorithms for sets and hash tables suffer from serious drawbacks that prevent ...

8 Partitioned first-level cache design for clustered microarchitectures

Paul Racunas, Yale N. Patt
June 2003 **ICS '03:** Proceedings of the 17th annual international conference on Supercomputing
Publisher: ACM

Full text available: pdf(191.74 KB) Additional Information: full citation, abstract, references, cited by, index terms

The high clock frequencies of modern superscalar processors make the wire delay incurred in moving data across the processor chip a significant concern. As frequencies continue to increase, it will become more difficult for a centralized first level ...

Keywords: clustered microarchitecture, partitioned cache

9 A low-cost memory remapping scheme for address bus protection

Lan Gao, Jun Yang, Marek Chrobak, Youtao Zhang, San Nguyen, Hsien-Hsin S. Lee
September 2006 **PACT '06:** Proceedings of the 15th international conference on Parallel architectures and compilation techniques
Publisher: ACM

Full text available: pdf(536.42 KB) Additional Information: full citation, abstract, references, index terms

The address sequence on the processor-memory bus can reveal abundant information about the control flow of a program. This can lead to critical information leakage such as encryption keys or proprietary algorithms. Addresses can be observed by attaching ...

Keywords: address bus leakage protection, secure processor

10 Substituting associative load queue with simple hash tables in out-of-order microprocessors

Alok Garg, Fernando Castro, Michael Huang, Daniel Chaver, Luis Piñuel, Manuel Prieto
October 2006 **ISLPED '06:** Proceedings of the 2006 international symposium on Low power electronics and design
Publisher: ACM

Full text available: pdf(174.77 KB) Additional Information: full citation, abstract, references, cited by, index terms

Buffering more in-flight instructions in an out-of-order microprocessor is a straightforward and effective method to help tolerate the long latencies generally associated with off-chip memory accesses. One of the main challenges of buffering a large ...

Keywords: LSQ, hash table, memory disambiguation, scalability

11 Using Address Independent Seed Encryption and Bonsai Merkle Trees to Make Secure Processors OS- and Performance-Friendly

Brian Rogers, Siddhartha Chhabra, Milos Prvulovic, Yan Solihin
December 2007 **MICRO '07:** Proceedings of the 40th Annual IEEE/ACM International Symposium on Microarchitecture
Publisher: IEEE Computer Society

Full text available: pdf(954.09 KB) Additional Information: full citation, abstract

In today's digital world, computer security issues have become increasingly important. In particular, researchers have proposed designs for secure processors which utilize hardware-based mem- ory encryption and integrity verification to protect the privacy ...

12 Improving hash join performance through prefetching

Shimin Chen, Anastassia Ailamaki, Phillip B. Gibbons, Todd C. Mowry
August 2007 **ACM Transactions on Database Systems (TODS)**, Volume 32 Issue 3
Publisher: ACM

Full text available: pdf(904.25 KB) Additional Information: full citation, abstract, references, index terms

Hash join algorithms suffer from extensive CPU cache stalls. This article shows that the standard hash join algorithm for disk-oriented databases (i.e. GRACE) spends over 80% of its user time stalled on CPU cache misses, and explores the use of ...

Keywords: CPU cache performance, CPU cache prefetching, Hash join, group prefetching, software-pipelined prefetching

13 A Peer-to-Peer Replica Location Service Based on a Distributed Hash Table

Min Cai, Ann Chervenak, Martin Frank
November 2004 **SC '04:** Proceedings of the 2004 ACM/IEEE conference on Supercomputing
Publisher: IEEE Computer Society

Full text available: pdf(343.25 KB) Additional Information: full citation, abstract, references, cited by

A Replica Location Service (RLS) allows registration and discovery of data replicas. In earlier work, we proposed an RLS framework and described the performance and scalability of an RLS implementation in Globus Toolkit Version 3.0. In this paper, we ...

Keywords: Algorithms, Experimentation, Grid, Peer-to-Peer, Replication

14 Correlated load-address predictors

Michael Bekerman, Stephan Jourdan, Ronny Ronen, Gilad Kirshenboim, Lihu Rappoport, Adi Yoaz, Uri Weiser
May 1999 **ACM SIGARCH Computer Architecture News**, Volume 27 Issue 2
Publisher: ACM

Full text available: pdf(149.18 KB) Publisher Site Additional Information: full citation, abstract, references, cited by, index terms

As microprocessors become faster, the relative performance cost of memory accesses increases. Bigger and faster caches significantly reduce the absolute load-to-use time delay. However, increase in processor operational frequencies impairs the relative ...

Keywords: context-based predictor, global correlation, load-address prediction, predictor implementation, recursive data structures

15 Register Packing: Exploiting Narrow-Width Operands for Reducing Register File Pressure
Oguz Ergin, Deniz Balkan, Kanad Ghose, Dmitry Ponomarev
December 2004 **MICRO 37:** Proceedings of the 37th annual IEEE/ACM International Symposium on Microarchitecture
Publisher: IEEE Computer Society
Full text available: pdf(224.06 KB) Additional Information: full citation, abstract, cited by

A large percentage of computed results have fewer significant bits compared to the full width of a register. We exploit this fact to pack multiple results into a single physical register to reduce the pressure on the register file in a superscalar processor. ...

16 Correlated load-address predictors
Michael Bekerman, Stephan Jourdan, Ronny Ronen, Gilad Kirshenboim, Lihu Rappoport, Adi Yoaz, Uri Weiser
May 1999 **ISCA '99:** Proceedings of the 26th annual international symposium on Computer architecture
Publisher: IEEE Computer Society
Full text available: pdf(149.18 KB) Publisher Site Additional Information: full citation, abstract, references, cited by, index terms

As microprocessors become faster, the relative performance cost of memory accesses increases. Bigger and faster caches significantly reduce the absolute load-to-use time delay. However, increase in processor operational frequencies impairs the relative ...

Keywords: context-based predictor, global correlation, load-address prediction, predictor implementation, recursive data structures

17 Simultaneous reference allocation in code generation for dual data memory bank ASIPs
Ashok Sudarsanam, Sharad Malik
April 2000 **ACM Transactions on Design Automation of Electronic Systems (TODAES)**, Volume 5 Issue 2
Publisher: ACM
Full text available: pdf(156.30 KB) Additional Information: full citation, abstract, references, cited by, index terms

We address the problem of code generation for DSP systems on a chip. In such systems, the amount of silicon devoted of program ROM is limited, so application software must be sufficiently dense. Additionally, the software must be written so as to meet ...

Keywords: code generation, code optimization, graph labelling, memory bank assignment, register allocation

18 Mercury: supporting scalable multi-attribute range queries
Ashwin R. Bharambe, Mukesh Agrawal, Srinivasan Seshan
August 2004 **ACM SIGCOMM Computer Communication Review**, Volume 34 Issue 4
Publisher: ACM
Additional Information:

Full text available: pdf(1.29 MB) full citation, abstract, references, cited by, index terms

This paper presents the design of Mercury, a scalable protocol for supporting multi-attribute range-based searches. Mercury differs from previous range-based query systems in that it supports *multiple attributes* as well as performs *explicit* ...

***Keywords**: distributed hash tables, load balancing, peer-to-peer systems, random sampling, range queries*

19 A privacy enhancing mechanism based on pseudonyms for identity protection in location-based services

Oliver Jorns, Gerald Quirchmayr, Oliver Jung

January 2007 **ACSW '07: Proceedings of the fifth Australasian symposium on ACSW frontiers - Volume 68**, Volume 68

Publisher: Australian Computer Society, Inc.

Full text available: pdf(333.09 KB) Additional Information: full citation, abstract, references

Over the past years Mobile Business has gained significant progress not only because of higher transfer rates as well as advanced processing power and memory capabilities of networks and mobile devices but also because of novel location-based mobile ...

Keywords: hash function, identity management, location-based services, privacy, pseudonyms, security, telecommunication services

20 Cryptography as an operating system service: A case study

Angelos D. Keromytis, Jason L. Wright, Theo De Raadt, Matthew Burnside

February 2006 **ACM Transactions on Computer Systems (TOCS)**, Volume 24 Issue 1

Publisher: ACM

Full text available: pdf(669.12 KB) Additional Information: full citation, abstract, references, index terms, review

Cryptographic transformations are a fundamental building block in many security applications and protocols. To improve performance, several vendors market hardware accelerator cards. However, until now no operating system provided a mechanism that allowed ...

Keywords: Encryption, authentication, cryptographic protocols, digital signatures, hash functions

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2008 ACM, Inc.
Terms of Usage Privacy Policy Code of Ethics Contact Us

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player

IEEE Xplore® RELEASE 2.5

Home | Login | Logout | Access Information | Alerts | Purchase History | Cart |

Welcome United States Patent and Trademark Office

Search Results BROWSE SEARCH IEEE XPLORE GUIDE SUP

Results for "(((memory)<in>metadata) <and> ((mac address)<in>metadata))<and> ((hash)&..." e-mail
Your search matched **3** of **1752210** documents.
A maximum of **100** results are displayed, **25** to a page, sorted by **Relevance** in **Descending** order.

New [Beta] Application Notes POWERED BY SPEC

Modify Search

(((memory)<in>metadata) <and> ((mac address)<in>metadata))<and> ((hash)<in> Search

Check to search only within this results set

Display Format: Citation Citation & Abstract

» Search Options

View Session History

New Search

» Key

IEEE JNL	IEEE Journal or Magazine
IET JNL	IET Journal or Magazine
IEEE CNF	IEEE Conference Proceeding
IET CNF	IET Conference Proceeding
IEEE STD	IEEE Standard

IEEE/IET Books Educational Courses Applicatio

IEEE/IET journals, transactions, letters, magazines, conference proceedings, and standard

view selected items Select All Deselect All

1. **Effect of Hash Collisions on the Performance of LAN Switching Devices and Netw**
Huntley, C.; Antonova, G.; Guinand, P.;
Local Computer Networks, Proceedings 2006 31st IEEE Conference on
Nov. 2006 Page(s):280 - 284
Digital Object Identifier 10.1109/LCN.2006.322112
AbstractPlus | Full Text: PDF(177 KB) IEEE CNF
Rights and Permissions

2. **A Hardware-Engine for Layer-2 classification in low-storage, ultra high bandwidth**
Papaefstathiou, V.; Papaefstathiou, I.;
Design, Automation and Test in Europe, 2006. DATE '06. Proceedings
Volume 2, 6-10 March 2006 Page(s):1 - 6
AbstractPlus | Full Text: PDF(296 KB) IEEE CNF
Rights and Permissions

3. **A memory efficient, 100 Gb/sec MAC classification engine**
Papaefstathiou, V.; Papaefstathiou, I.;
Local Computer Networks, 2005. 30th Anniversary. The IEEE Conference on
15-17 Nov. 2005 Page(s):2 pp.
Digital Object Identifier 10.1109/LCN.2005.9
AbstractPlus | Full Text: PDF(144 KB) IEEE CNF
Rights and Permissions

Help Contact Us Privacy &

Indexed by Inspec®

© Copyright 2008 IEEE –

Google

hash table and memory MAC address Search Advanced Search Preferences

The "AND" operator is unnecessary -- we include all search terms by default. [details]

Web Results 1 - 10 of about 630,000 for hash table and memory MAC address. (0.27 seconds)

Forwarding table incorporating hash table and content addressable ...
A hash table is realized using a large enough memory to store the hashed index ... The data stored in the hash table is limited to one MAC address per entry ...
www.patentstorm.us/patents/6735670-description.html - 49k - Cached - Similar pages

OpenWrtDocs/Hardware/Huawei/EchoLife HG520 - OpenWrt
PID hash table entries: 64 (order 6: 512 bytes) Using 120.000 MHz high ... 2006 19:39:17 usb0: MAC Address: 00 11 F5 E6 9D 00 usb0: Host MAC Address: 00 11 ...
wiki.openwrt.org/OpenWrtDocs/Hardware/Huawei/EchoLife_HG520 - 28k - Cached - Similar pages

A Memory Efficient, 100 Gb/sec MAC Classification Engine
(HBCE) compacts the MAC address tables and. supports high speed decisions using significantly less. memory than the existing solutions. 2. Hash Based ...
ieeexplore.ieee.org/iel5/10397/33047/01550891.pdf?isnumber=33047&prod=CNF&arnumber=1550891&ar... - Similar pages

Linux on the Blackfin processor > Projects > Das U-Boot for the ...
Freeing unused kernel memory: 56k freed (0x12b000 - 0x138000) ... FLASH: 4 MB In: serial Out: serial Err: serial Using MAC Address 92:65:09:D1:3C:71 Bus 0: ...
blackfin.uclinux.org/gf/project/u-boot/tracker/?action=TrackerItemEdit&tracker_item_id=1735 - 36k - Cached - Similar pages

Computer-based system for validating hash-based table lookup ...
6438674, Perloff, 711/216, Hash Cam having a reduced size memory array and its ..... In particular, a user may program the address table 44 with MAC and IP ...
www.freepatentsonline.com/6757742.html - 52k - Cached - Similar pages

Samsung SWL-2100P no MAC address
2392.06 BogoMIPS Memory: 256960k/262064k available (974k kernel code, 4716k reserved, 395k data, 84k init, 0k highmem) Dentry cache hash table entries: ...
lists.shmoo.com/pipermail/hostap/2003-March/001898.html - 13k - Cached - Similar pages

Mac OS X Manual Page For Hash(3tcl)
Tcl_DeleteHashTable deletes all of the entries in a hash table and frees up the memory associated with the table's bucket array and entries. ...
developer.apple.com/documentation/Darwin/Reference/ManPages/man3/Hash.3tcl.html - 26k - Cached - Similar pages

BT Voyager 220V Serial Port
... of MAC Addresses (1-32) : 3 Ethernet MAC Address : 00:11:f5:cf:1b:0a Memory ... TCP, IGMP IP: routing cache hash table of 512 buckets, 4Kbytes TCP: Hash ...
jcsu.jesus.cam.ac.uk/~acw43/projects/voyager/serial.html - 8k - Cached - Similar pages

ARM9 SBC
... 32768 bytes) Inode-cache hash table entries: 4096 (order: 2, 16384 bytes) Memory: 32MB ... Configuring network interfaces...eth0: Setting MAC address to ...
www.ime.usp.br/~fr/sbc/ - 14k - Cached - Similar pages

Amitesh Singh - Can I change 'RedBoot Config' using MTD driver. ?
... bytes) Inode-cache hash table entries: 8192 (order: 3, 32768 bytes) Memory: .... OK +

Panther Rev 1.2 Redboot Release 1.1 Ethernet eth0: **MAC address ...**
sourceware.org/ml/ecos-discuss/2006-09/msg00002.html - 15k - Cached - Similar pages

1 2 3 4 5 6 7 8 9 10 **Next**

---

hash table and memory MAC addres Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve | Try Google Experimental

---

©2008 Google - Google Home - Advertising Programs - Business Solutions - About Google

Google

hash table and memory MAC address and CA Search Advanced Search Preferences

The **"AND"** operator is unnecessary -- we include all search terms by default. [details]

Web Results **1** - **10** of about **162,000** for **hash table and memory MAC address and CAM**. (**0.26** seconds)

Forwarding **table** incorporating **hash table** and content addressable ...
The data stored in the **hash table** is limited to one **MAC address** per entry thus ... **memory** (**CAM**), adding a **MAC address** entry to the forwarding **table** by ...
www.patentstorm.us/patents/6735670-description.html - 49k - Cached - Similar pages

> Forwarding **table** incorporating **hash table** and content addressable ...
> The **CAM** is used when an **address** cannot be found in the **hash table**. When **MAC** addresses are being added to the forwarding **table**, they are first tried in the ...
> www.patentstorm.us/patents/6735670.html - 15k - Cached - Similar pages
> More results from www.patentstorm.us »

Combined **hash table** and **CAM address** recognition in a network ...
Combined **hash table** and **CAM address** recognition in a network. Document Type and Number:. European Patent EP0522743. Kind Code:. B1. Link to this page: ...
www.freepatentsonline.com/EP0522743.html - 63k - Cached - Similar pages

Content-addressable **memory** - Wikipedia, the free encyclopedia
Content-addressable **memory** (**CAM**) is a special type of computer **memory** used in ... The **MAC address table** is usually implemented with a binary **CAM** so the ...
en.wikipedia.org/wiki/Content-**address**able_**memory** - 28k - Cached - Similar pages

The ARP cache structure of lan emulation server in ATM-LAN system ...
address. For a given **hash table** with m cache entries with. total of n **MAC**-ATM binding information, ... Comparand was found in **memory**. Thus, with a **CAM**, ...
ieeexplore.ieee.org/iel5/6630/17685/00818389.pdf - Similar pages

OpenWrt / ASUS WL-600g Serial connection log
Board IP **address** : 192 Inode-cache **hash table** entries: 2048 (order: 1, 8192 bytes) Host IP **address** : 192.168.1.100 **Memory**: 13340k/16000k available (1578k ...
forum.openwrt.org/viewtopic.php?pid=62517 - 23k - Cached - Similar pages

EZchip Technologies
DRAM offers 28 times the density of the largest **CAM** or SRAM (512 Mbit versus .... VPLS **tables** to **memory** mapping. The MPLS tunnel **hash table** and the **MAC hash** ...
www.ezchip.com/t_**memory**_whpaper.htm - 29k - Cached - Similar pages

Linux on the Blackfin processor > Projects > uClinux for the ...
Using **MAC Address** FF:FF:FF:FF:FF:FF TFTP from server 10.100.4.174; our IP **address** is 10.100. .... TCP: **Hash tables** configured (established 2048 bind 2048) ...
blackfin.uclinux.org/gf/project/uclinux-dist/mailman/?action=ListThreads&mailman_id=41&_forum_act... - 31k - Cached - Similar pages

BT Voyager 220V Serial Port
210.53 BogoMIPS **Memory**: 14140k/16000k available (1148k kernel code, 1860k reserved, 84k data, 4) Dentry-cache **hash table** entries: 2048 (order: 2, ...
jcsu.jesus.**cam**.ac.uk/~acw43/projects/voyager/serial.html - 8k - Cached - Similar pages

Details and datasheet on part: RTL8316

Address look-up **table** consists of 8K entries of **hash table** and a 128 entries of **CAM**. The RTL8308B uses **address hashing** algorithm or direct mapping method to ...
www.digchip.com/datasheets/parts/datasheet/398/RTL8316.php - 17k -

Cached - Similar pages

1 2 3 4 5 6 7 8 9 10 **Next**

---

hash table and memory MAC addres Search

Search within results | Language Tools | Search Tips | Dissatisfied? Help us improve | Try Google Experimental

---

©2008 Google - Google Home - Advertising Programs - Business Solutions - About Google